

Customer No. 20350

TOWNSEND and TOWNSEND and CREW LLP
Two Embarcadero Center, 8th Floor
San Francisco, California 94111-3834
(415) 576-0200

ASSISTANT COMMISSIONER FOR PATENTS
BOX PATENT APPLICATION
Washington, D.C. 20231



Sir:

Transmitted herewith for filing under 37 CFR 1.53(b) is the

- patent application of
 continuation patent application of
 divisional patent application of
 continuation-in-part patent application of

Inventor(s)/Applicant Identifier: David P. Rossum

For: SAMPLE RATE CONVERTER HAVING DISTRIBUTED FILTERING

Enclosed are:

- 24 page(s) of specification
 5 page(s) of claims
 1 page of Abstract
 12 sheet(s) of [] formal informal drawing(s).
 An assignment of the invention to Creative Technology Ltd.
 A signed Declaration & Power of Attorney
 Assignment and Cover Sheet

	(Col. 1)	(Col. 2)
FOR:	NO. FILED	NO. EXTRA
BASIC FEE		
TOTAL	33	-20
CLAIMS	=	*13
INDEP.	4	-3
CLAIMS	=	*1
[] MULTIPLE DEPENDENT CLAIM PRESENTED		

* If the difference in Col. 1 is less than 0, enter "0" in
Col. 2.

Please charge Deposit Account No. 20-1430 as follows:

- Filing fee \$ 1,072.00
 Any additional fees associated with this paper or during the pendency of this application.
 The issue fee set in 37 CFR 1.18 at or before mailing of the Notice of Allowance, pursuant to 37 CFR 1.311(b)

- A check for \$ _____ is enclosed.
2 extra copies of this sheet are enclosed.

Telephone:
(415) 576-0200

Faxsimile:
(415) 576-0300

Attorney Docket No. 17002-014600US

"Express Mail" Label No. EL380683423US

Date of Deposit: October 27, 1999

I hereby certify that this is being deposited with the United States
Postal Service "Express Mail Post Office to Addressee" service
under 37 CFR 1.10 on the date indicated above, addressed to:

Assistant Commissioner for Patents
Washington, D.C. 20231

By:

Dave Shissler



SMALL ENTITY

RATE	FEES
	\$380.00
x \$9.00 =	
x \$39.00 =	
+ \$130.00 =	
TOTAL	

OTHER THAN
SMALL ENTITY

OR RATE	FEES
OR	\$760.00
OR	x \$18.00 = \$234.00
OR	x \$78.00 = \$78.00
OR	+ \$260.00 =
OR	TOTAL \$1,072.00

Respectfully submitted,

TOWNSEND and TOWNSEND and CREW LLP

Paul C. Haughey
Paul C. Haughey
Reg No.: 31,836
Attorneys for Applicant

PATENT APPLICATION
SAMPLE RATE CONVERTER HAVING DISTRIBUTED FILTERING

Inventor:

David P. Rossum,
a citizen of United States, residing at,
25513 Paseo de Cumbre
Monterey, CA, 93940
USA

Assignee:

Creative Technology Ltd.
31 International Business Park
Creative Resource
Singapore, Singapore 609921

Entity: Large

SAMPLE RATE CONVERTER HAVING DISTRIBUTED FILTERING

BACKGROUND OF THE INVENTION

In many electronic applications, signals are represented and processed 5 digitally. Digital words, or samples, represent the value of the signal at a regular time interval. This regular interval is often referred to as the sample rate, and is typically expressed in units of kilohertz (kHz) representing the reciprocal of the sample interval time period. The signal thus represented can have no energy above half the sample rate; the frequency equal to half the sample rate is called the Nyquist frequency.

Practically, the signal thus represented must have its energy limited at some 10 frequency below the Nyquist frequency. The band of frequencies which the signal is intended to contain is called the passband of the signal, and the upper frequency of the passband is called the passband edge. The band of frequencies between the passband edge and the Nyquist frequency is called the guardband. For example, for audio signals sampled at 15 48 kHz, the Nyquist frequency is 24 kHz, the passband edge is generally defined as 20 kHz, and the width of the guardband is 4 kHz.

There are situations when the available sample rate of the data is different 20 from the desired sample rate. Depending on the characteristics of the sample data and how much the available and desired rates differ, several approaches may be used to convert the signal at one sample rate to a signal at another sample rate without substantially altering the meaning of the signal. A first common technique for sample rate reduction is called decimation. This technique is employed to reduce the sample rate of a signal by an integer divisor d. To reduce aliasing effects, any frequency components in the signal above the 25 passband edge of the output sample rate are typically attenuated by applying a low pass filter to the incoming signal before reducing the sample rate. The rate reduction is performed by simply periodically discarding $d - 1$ input samples of the signal, thus causing the output to consist of every d^{th} filtered input sample.

A second common technique for sample rate conversion is known as up-sampling. This technique results in an increase of the sample rate of a signal by an integer 30 factor u. The increased sample rate is achieved by augmenting the number of samples that represent the signal by inserting $u - 1$ zero value samples between adjacent input samples. The resulting samples are typically filtered through an anti-imaging low pass filter at the

higher output sample rate to remove frequency components above the original Nyquist frequency.

Typically, the decimation and up-sampling techniques are combined in a multi-stage converter to obtain a desired sample rate. The “classical” (prior art) multi-stage converter algorithm involves successive zero insertion, filtering, and decimating of a signal. For example, to change the sample rate of a signal sampled at rate R to a new rate $R' = R*L/M$, a number of zero value samples equal to $L-1$ are inserted between each sample of the signal (L and M are integers which can slowly change with time). This creates a signal at sample rate $L*R$. Lowpass filtering of the signal is applied that removes all frequencies above the lower of the Nyquist frequency associated with the original signal or the new Nyquist frequency associated with the signal at the converted sample rate. Then, the signal is decimated by the factor M by deleting all but every M^{th} sample, producing a new signal at rate R' .

The computational complexity of the aforementioned sample rate conversion techniques is a function of the filter employed, which when properly designed depends on the width of the guardband. The guardband width should determine the width of the filter’s transition band, which is the band of frequencies between the passband edge and the lower edge of the filter’s stopband. The lower edge of the lowpass filter’s stopband is the frequency above which all frequencies are attenuated at least as much as the stopband attenuation specification. Simply stated, the higher the quality of the filtering and the more narrow the guardband, the more computationally intensive the converter. Important characteristics to consider when determining the quality of a filter include stop-band attenuation, passband ripple, and numerical accuracy. Passband ripple in a sample rate converter results from variation in gain of the filter that is expressed as a worst case deviation in decibels from nominal gain. A high quality sample rate converter for audio will generally have a passband ripple less than ± 0.01 dB. While such values of ripple are far below the limit of audibility, they are important if the signal may be passed through a sample rate converter multiple times, as might occur in a recording studio.

The level of distortion in a sample rate converter is a function of the magnitude of aliases typically produced during decimation. These result from the non-ideal stop-band attenuation of the filter in any real sample rate converter, which will cause some aliases to be produced by the decimation step. Aliasing may be measured by testing the converter with a full level sine waveform of varying frequency, noting the peak magnitude of

any undesired components using the standard “peak THD+N” measurement technique. The ratio of the full level signal to the THD+N expressed in decibels, is the alias rejection of the converter at that signal frequency, which is directly related to the stopband attenuation of the filter and to the distortion of the converter. A high quality sample rate converter will have an
5 alias rejection of more than 96 dB throughout the audio spectrum. It is worthy of note that the distortion caused by inadequate alias rejection is proportional to signal level. This is in marked contrast to quantization distortion or dither noise, both of which remain at a constant level regardless of signal level. Numerical accuracy of a filter depends intimately on the design details of the filter. In sample rate converters, the filter is generally designed so that
10 the distortion resulting from numerical inaccuracy is substantially below the distortion generated by aliasing.

Classical algorithm single stage interpolation when converting to equal or higher sampling rate with equal or wider passband.

Figs. 1A, 1B and 1C illustrate filter requirements for a prior art classical algorithm single stage interpolator when converting a signal to an approximately equal or higher sampling rate with an equal or wider passband, while meeting certain quality specifications. Fig. 1A shows the spectrum of the input signal at sample rate F_{sx} , including passband frequency range 101, guardband 102, first image 103, second image 104, third image 105, and fourth image 106. The images are redundant copies of the original passband frequency range whose production is inherent in the process of digital sampling. Images above the fourth are not shown. Fig. 1B shows the requirements for the filter of a single stage classical interpolator, including passband ripple 111, narrow transition band 112 with width less than twice the width of the guardband, beginning no lower than the upper edge of input signal passband 101 and ending below the bottom edge of the first passband image 103,
15 and stopband 113 with finite attenuation. Fig. 1C shows the result of the interpolator filtering operation with passband 121 unchanged (apart from the effects of passband ripple) and all images attenuated by the stopband, thus preventing any significant aliasing when decimation occurs in the final interpolator operation.

One fundamental decision regarding performance of sample rate converters
30 involves the interpretation of the guardband, the frequencies between the top of the passband and the Nyquist frequency. In particular, the choice must be made as to whether the filter stopband should begin at the Nyquist frequency, or at the first image of the passband edge. The latter case is less conservative, but requires approximately a factor of two less computation. The audible consequences of this choice are fairly subtle. If the less

conservative choice is made, any energy present in the incoming signal above the input passband will result in aliases greater than the stopband limit, while the more conservative approach will not produce aliases under any conditions. When the conversion is to a lower sample rate, the additional aliases will always lie above the passband of the output sample
5 rate. When the conversion is to a higher rate, they will lie well above the passband of the incoming rate. In both cases, if the lower of the passbands is assumed to be the limit of hearing, the aliases will be inaudible and hence of no consequence, which is discussed below with respect to Figs. 2A-2G and 3A-3H.

Output sample rate greater than input sample rate.

Fig. 2A shows the spectrum of the input signal at sample rate F_{SX} , including passband frequency range 201, input guardband 202, first image 203, second image 204, third image 205, and fourth image 206. Images above the fourth are not shown. Also shown is a signal component 207 within the guardband, and that signal's first image 208, and higher images 209. Fig. 2B shows a less conservative interpolator filter having a transition band 212 of twice the width of the input guardband. Fig. 2C shows the result of the interpolator filtering operation shown in Fig. 2B with signal 207 and image 208 both partially attenuated as filtered signals 227 and 228 respectively. Fig. 2D shows the result of the decimation operation on the signal of 2C when the output sample rate F_{SY} is above the input rate, showing the results of signal 207 and image 208 above the input passband as signal components 237 and 238. Fig. 2E shows a more conservative interpolator filter having a transition band 242 of the same width as that of the guardband. Fig. 2F shows the result of the interpolator filtering operation shown in Fig. 2E with signal 207 partially attenuated as filtered signal 257, but image 208 completely attenuated to the stopband limit. Fig. 2G shows the result of the decimation operation on the signal of 2F when the output sample rate F_{SY} is above the input rate, showing the results of signal 207 above the input passband as signal component 267, but notably with image 208 absent.

Output sample rate less than input sample rate.

Fig. 3A shows the spectrum of the input signal at sample rate F_{SX} , including passband frequency range 301, input guardband 302, first image 303, second image 304, third image 305, and fourth image 306. Images above the fourth are not shown. Fig. 3B shows the output sample rate F_{SY} which is lower than the input rate F_{SX} and consequently has smaller output passband 311 and output guardband 312. Fig. 3C shows a less conservative interpolator filter having a transition band 322 of twice the width of the output guardband 312. Fig. 3D shows the result of the interpolator filtering operation of 3C with a portion of

the input passband 301 partially attenuated as filtered frequency range 337. Fig. 3E shows the result of the decimation operation on the signal shown in Fig. 3D when the output sample rate F_{sy} is below the input rate, showing the results of signals in frequency range 337 as both signals 347 and aliases 348, both of which lie within the frequencies of the output guardband 5 312. Fig. 3F shows a more conservative interpolator filter having a transition band 352 of the same width as the output guardband 312. Fig. 3G shows the result of the interpolator filtering operation of 3F with a narrower portion of the input passband 301 partially attenuated as filtered frequency range 367. Fig. 3H shows the result of the decimation operation on the signal of 3G when the output sample rate F_{sy} is below the input rate, showing the results of 10 signals in frequency range 367 as signals 377 which lie within the frequencies of the output guardband 312, but with all aliases absent.

The performance measurement consequences of different guardband interpretations are more definite. If the converter is tested with input frequencies only below the lower of the passband limits, then either interpretation will meet the design specifications. However, if a converter going to a higher sample rate is tested with inputs containing energy in frequencies above the input passband limit, the aliases of these frequencies will exceed the stopband limit. If a converter going to a lower sample rate is tested with input frequencies above the output passband limit, and measured without a brick-wall filter at the output passband limit, then aliases above the stopband will be measured. In general, testing is limited to within the passbands, so the difference between the interpretations is generally not noted.

In the context of this disclosure, unless explicitly stated otherwise, it will be assumed that the guardband is interpreted in the less conservative manner, in which guardband aliases are acceptable. Thus for the case of the “classical” sample rate converter 25 algorithm explained above, the filter is designed such that the transition band extends from the edge of the passband to the first image of the edge of the passband, thus having a width twice that of the guardband. In other words, for a comparable prior art “high quality” sample rate converter operating on audio at a sample rate of 48 kHz, the filter specifications would include passband ripple less than ± 0.01 dB from 0 to 20 kHz, stopband rejection greater than 30 96 dB above 28 kHz, and a transition band of 8 kHz width from 20 to 28 kHz, as previously described for Fig. 1.

Different types of sample rate converters.

Figs. 4A-4E illustrates the frequency domain behavior of a number of prior art, common, low complexity, sample rate converter techniques. Fig. 4A shows the spectrum of the input signal for any of these converters at a sample rate F_{SX} including passband frequency range 401, input guardband 402, first image 403, second image 404, third image 405, and fourth image 406. Images above the fourth are not shown.

The simplest sample rate converter is called a “drop sample” interpolator, and uses a filter with a one sample wide rectangular impulse, which can be implemented with no computation at all. The complexity is very low because no arithmetic operations are required, and the quality is very poor. The frequency response of the filter associated with this sample rate converter is found by taking the magnitude of the Fourier transform of its impulse response, which is $\sin^2(\pi f/f_s) / (\pi f/f_s)$, where f is the frequency and f_s is the sample rate. This frequency response is shown in Fig. 4B, and includes passband ripple 411 of magnitude approximately 2.6 dB, stopband rejection 412 of magnitude 13dB, and transition band 413 which extends (at a sample rate of 48 kHz) from the 20 kHz input signal passband edge to 39 kHz, thus having a width of 19 kHz.

The next simplest sample rate converter is a linear interpolator. The complexity is one multiply and two adds per output sample, and the quality is considerably better than drop sample. The impulse response of the associated filter, a triangular function two samples in width, has Fourier transform $\sin^2(\pi f/f_s) / (\pi f/f_s)^2$. The associated frequency response is shown in Fig. 4C, and includes passband ripple 421 of magnitude approximately 5.3 dB, stopband rejection 422 of magnitude 27dB, and transition band 423 which extends (at a sample rate of 48 kHz) from the 20 kHz input signal passband edge to 39 kHz, thus having a width of 19 kHz.

These two classical algorithm sample rate converters are generally classified as low quality and low complexity. While quantitative distortion measurements of these two converters can be made, the results for broadband signals are so poor as to be nearly meaningless. A linear interpolator works fairly well when the input signal energy is concentrated at very low frequencies (<1/10 of the Nyquist frequency).

Sample rate converters using the classical algorithm can be constructed from higher order mathematical interpolation techniques related to the drop sample and linear interpolators. One such family of interpolators are the splines, also called B-splines or (for the one of third order) cubic splines. The Fourier transform of the Nth order spline

interpolator has been shown to be $\sin^n(\pi f/f_s)/(\pi f/f_s)^n$, sometimes abbreviated sinc^n . Fig. 4D shows the frequency response of several members (orders 3, 4, 7, 15 and 31) of the spline family, including passband ripples 431, stopbands 432, and transition bands 433 which in all cases extend (at a sample rate of 48 kHz) from the 20 kHz input signal passband edge to approximately 39 kHz, thus all having a width of 19 kHz. It is interesting to note that the transition band width of a spline is fixed and approximately independent of its order, and also to note that while the stopband rejection of higher order splines is quite good, the passband ripple of a spline becomes increasingly poorer with increased order. Thus, high order splines are typically not useful for high quality sample rate conversion.

Another family of interpolators are polynomial interpolators, which are generally implemented according to the method of Lagrange, hence called Lagrangian interpolators. A closed form of the frequency response of the Nth order Lagrangian interpolator is too complex to reproduce here, but Fig. 4E shows the frequency response of several members of this family (orders 3, 4, 7, 15 and 31). Note passband ripples 441, stopbands 442, and transition bands 443. It is noteworthy that while the transition bands of higher order Lagrangian interpolators become narrower with increasing order, the stopband rejection remains quite poor. Thus higher order Lagrangian interpolators are not useful for high quality sample rate conversion.

Drop sample, linear, spline and Lagrangian interpolators have also been used in multistage systems for sample rate conversion. Because of the above mentioned limitations in filter quality for such interpolators, a high degree of oversampling (typically 128 times) must be used to achieve high quality. The computational complexity of the required oversampling is a major drawback to this approach.

Sample rate converters using the classical algorithm can also be constructed using FIR (Finite Impulse Response) filters. For converters where the sample rate ratio is fixed, polyphase FIR filters are generally used. For variable sample rate ratios, an FIR filter impulse response is typically stored in a table and interpolated, most commonly using linear interpolation.

The FIR filters used in such sample rate converters are generally designed to meet the requirements illustrated in Figs. 1A-1C. When a windowed sinc function is used, the width of the transition band and the filter quality can be independently controlled, but in general there is no independent control of stopband rejection and passband ripple. When optimization methods, such as linear programming or Remez exchange are used, independent

and precise control of passband ripple, stopband rejection, and transition band width can be accomplished, even to the extent of further dividing the bands to produce filters with more precisely controlled frequency responses.

Intermediate quality sample rate converters can be constructed with FIR filter
5 orders from four to sixteen. These converters can be specified for distortion of wideband signals, although in general the results are substantially inferior to distortion measurements for other digital audio subsystems. An Nth order FIR sample rate converter, using linear interpolation of the FIR filter coefficients, will have a computational complexity of $2N$ multiplies (one for the convolution and one for the linear interpolation) and $3N$ additions (one
10 for the convolution and two for the linear interpolation) per output sample.

High quality sample rate converters attempt to equal the performance of typical high fidelity digital audio subsystems. They typically involve FIR filters with orders from thirty-two to several hundred.

The filter associated with a sample rate converter may be either an infinite impulse response (IIR) filter or a finite impulse response (FIR) filter. Applying an IIR filter in the "classical" algorithm, is usually useful when efficiency is not of great concern and R/R' (sample rate R/new sample rate R') is a ratio of small integers. The primary disadvantage of the method is that the IIR filter is only efficient with a low rate LR , because the IIR filter is recursive. Thus, each output sample from the IIR filter depends
20 computationally on the previous outputs, and thus must be computed at the LR rate. Also, it is not practical to vary R/R' in real time because the restriction to small integer ratios (allowing LR to be small enough to be realizable) causes changes in rate to abruptly alter coefficients, producing audible defects.

Applying an FIR filter in the "classical" algorithm is much more practical than
25 the IIR filter, because the FIR filter is not recursive. Output samples which are dropped during decimation need not be calculated. Similarly, multiplicative operations upon inserted zeroes in the incoming sample stream need never be performed. This implies that the computational complexity of this approach is independent of the value of L . In other words, a fixed number of multiply-add operations must be performed for each output sample. In
30 general, despite the poorer efficiency of FIR filters in terms of computational steps to produce a desired filter specification, this approach is superior to the IIR approach except possibly when R/R' is a small integer ratio.

As can be seen, the main drawback of high quality sample rate converters involves the mathematical complexity that typically results from providing high quality filtering.

What is needed, therefore, is a sample rate conversion technique of high quality and reduced computational complexity.

SUMMARY OF THE INVENTION

The present invention is a method and a computer program product for sample rate conversion that features distributive (or hybrid) filtering to minimize unwanted artifacts such as aliasing, and which has low computational requirements while avoiding the aforementioned artifacts. The method includes receiving, at a first sample rate, a plurality of data points, associated with a first signal, operating on the plurality of data points to associate the signal with a predetermined set of parameters, with the set of parameters including a first transition band having an image corresponding thereto; and varying the sample rate associated with the first signal by interpolation with an interpolator having associated therewith a second transition band, with the width associated with the second transition band being a function of a spectral separation between the first transition band and its image, wherein a second signal is produced having a sequence of data samples approximating the first signal.

In this manner, the width of the second transition band may be relaxed. By relaxing the width of the second transition band, the computational requirements involved with varying the sample rate by interpolation are reduced, while the unwanted artifacts are substantially eliminated. To that end, varying the sample rate may produce output data samples associated with the second signal by convolving a predetermined finite number N of data points with an equal number of coefficients, with N being greater than two. The number N is determined by the desired converter quality and the allowable width of the second transition band, and the values of the coefficients are typically a function of the temporal spacing between the output data sample and the corresponding data points.

The present invention in one embodiment uniquely provides an interpolation filter with parameters that are independently adjustable. In particular, an FIR interpolator allows a wide transition band to be specified when the interpolator is programmed, while independently adjusting the passband ripple and stopband rejection to optimize the results.

The filtering is distributed or hybrid because, in a preferred embodiment, the interpolation filter is split into a halfband filter and a FIR or IIR interpolation filter. The

halfband filter is a simple filter that can easily eliminate the first and second images. The elimination of these images allows the FIR or IIR to have a broad transition band, making its design simpler. The particular placement of the halfband filter preferably is varied depending on the ratio of the desired output and input sample rates.

5 In a first embodiment, when the ratio R of the output sample rate divided by the input sample rate is greater than or approximately equal to one (e.g., 44kHz to 48kHz), the preferred embodiment includes up-sampling and filtering the plurality of data points by a factor of two before varying the sample rate. The halfband filter is included in the upsampling.

10 In a second embodiment, when the input sample rate is approximately twice the output sample rate (e.g., 96kHz to 48kHz), the halfband filter is used instead of upsampling, followed by the interpolation filter.

In a third embodiment, when the ratio of the output sample rate to the input sample rate is less than or approximately equal to one, yet significantly greater than $1/2^n$ (e.g., 48 kHz to 44kHz), the present invention includes filtering and decimating the plurality of data points after varying the sample rate.

In a fourth embodiment, when the ratio of the output sample rate to the input sample rate can be in a large range, and is only known to be larger than $1/(2^n)-1$ yet another embodiment of the present invention combines up-sampling and decimation, with a halfband filter being used on each of these steps. In between, the interpolation filter is used.

BRIEF DESCRIPTION OF THE DRAWINGS

Figs. 1A, 1B and 1C are graphical representations of filtering parameters for a prior art classical algorithm employed to achieve single stage interpolation when converting a signal to an approximately equal or higher sampling rate with an equal or wider passband;

25 Figs. 2A-2G are graphical representations showing operational characteristics of a prior art sample rate converter operating to provide an output sample rate that is greater than the input sample rate;

30 Figs. 3A-3H are graphical representations showing the operational characteristics of a prior art sample rate converter operating to provide an output sample rate that is less than the input sample rate;

Figs. 4A-4E are graphical representations showing the frequency domain behavior of various low complexity, prior art sample rate converter techniques;

Fig. 5 is a block diagram of a signal processing system suitable for implementing the present invention;

Fig. 6 is a block diagram of a sample rate converter, in accordance with the present invention;

5 Figs. 7A-7E are graphical representations of the parameters of an upsampling filter and interpolator in accordance with the present invention;

Fig. 8A is a flow diagram of power symmetric half-band upsampling elliptical filter is accordance with the present invention;

Fig. 8B is a flow diagram of typical allpass filters shown above in Fig. 8A;

10 Fig. 8C is a flow diagram for an IIR filter power symmetric half-band decimating elliptical filter in accordance with the present invention;

Fig. 9 is a simplified plan view of a compensating circuit where a first order allpass filter partially compensates phase nonlinearity of a power symmetric elliptical half-band filter, in accordance with the present invention;

15 Fig. 10 is a flowchart describing steps for converting a signal to differing sample rates using the processing system shown above in Fig. 5 to produce the sample rate converter shown in Fig. 6;

Fig. 11 is a block diagram of a sample rate converter, in accordance with an alternate embodiment of the present invention;

20 Figs. 12A-12G are graphical representations of the operational parameters of the half-band filter and interpolator shown above in Fig. 11;

Fig. 13 is a block diagram of a sample rate converter, in accordance with a second alternate embodiment of the present invention; and

25 Figs. 14A-14G are graphical representations of the operational parameters of the half-band filter and interpolator shown above in Fig. 13;

Fig. 15 is a block diagram of a sample rate converter, in accordance with a second alternate embodiment of the present invention; and

30 Figs. 16A-16G are graphical representations of the operational parameters of the half-band filter and interpolator shown above in Fig. 15.

DESCRIPTION OF THE SPECIFIC EMBODIMENTS

Overall System

Referring to Fig. 5, a signal processing system 500 suitable for implementing the present invention is shown including a host processor 502, RAM 504, ROM 506, an

interface controller 508, a display 510, a set of buttons 512, an analog-to-digital (A-D) converter 514, a digital-to-analog (D-A) converter 516, an interconnect block 518, a digital signal processor 520, a disk controller 522, a hard disk drive 524, and a floppy drive 526.

In operation, the signal processing system 500 captures samples of an analog signal, processes the samples, and outputs the processed samples. The present invention may be employed to process signals containing various types of information, such as audio information, control information and the like. To that end, A-D converter 514 converts analog signals to digital samples. Signal processing operations on the samples may be performed by host processor 502 or digital signal processor 520. Samples may be stored on hard disk drive 524 under the direction of disk controller 522. A user may request a particular signal processing operation using button set 512 and may view system status, or input or output waveforms on display 510. Once signals have been processed, they may be outputted using D-A converter 516 to convert samples to an analog signal.

The program control information for host processor 502 and DSP 520 is operably disposed in RAM 504. Long term storage of control information may be in ROM 506, on disk drive 524 or on a floppy disk 528 insertable in floppy drive 526. Interconnect block 518, which in one embodiment is an application-specific integrated circuit (ASIC), serves to interconnect and buffer between the various operational units. DSP 520 is preferably a 50 MHz TMS320C32 available from Texas Instruments. Host processor 502 is preferably a PentiumTM microprocessor available from Intel. Software to implement the present invention may be stored on a floppy disk 528, in ROM 506, on hard disk drive 524 or in RAM 504 at runtime. The ROM 506 stores information necessary to enable the signal processing system 500 to operate at differing sample rates while reducing the computational complexity to achieve the same. To that end, ROM 506 stores information that is operated on by host processor 502 or digital signal processor 520 to function as a sample rate converter capable of conversion by a rational, irrational or time varying ratio R. The ratio R is approximated by a rational value L/M, where L and M are integers that can slowly change with time.

It is also possible to design interconnect block 518 to perform the functions of DSP 520 in either a programmable or fixed algorithm form. Similarly, the DSP program can be contained in interconnect block 518, as can temporary information storage, if interconnect block 518 is an ASIC.

Sample Rate Converter

In a first embodiment, shown in Fig. 6, information stored in ROM 506 defines a sample rate converter 600 that converts an incoming stream of data samples x_n , representing an input signal, at a first sample rate to a second stream y_m , nominally at an equal, or higher, sample rate. This is achieved by having an upsampling filter 602 coupled to an interpolator 604 such that the data sample x_n are received by the filter 602. Although the interpolator 604 may consist of any known interpolation technique which can be parameterized for transition band width, it is preferred that interpolator 604 perform an Nth order FIR sum of products convolution on its input w_k to produce each output y_m such that

10

$$y_m = w_{k-N/2+1}C_0(f) + w_{k-N/2+2}C_1(f) + \dots + w_{k+N/2}C_{N-1}(f)$$

15
16
17
18
19
20

where y_m is the output sample interpolated at the fractional sample period f beyond sample k of signal w (i.e. $m = k+f$), and $C_i(f)$ is one of N coefficients which are computed by linearly interpolating from two values obtained from a table containing the impulse response of the FIR filter associated with the interpolator which has been designed to have the desired transition band width, passband ripple and stopband rejection characteristics.

An important consideration during sample rate conversion is to avoid aliasing. To that end, it is desirable to reduce, if not eliminate, energy corresponding to the input signal x_n above the Nyquist frequency associated with the incoming sample rate. Typically, however, satisfying this criterion increases the computational complexity of the sample rate converter 600. To reduce the computational complexity of the sample rate converter 600 while satisfying this criterion, the interpolator 604 is designed so that the transition band associated therewith has a width that is greater than twice the width of the guardband of the signal x_n received by the sample rate converter 600. In this manner, the computational complexity of the interpolator 604 may be relaxed. Relaxing the computational complexity of the interpolator 604 without sacrificing quality is achieved by over-sampling the input signal x_n . Although the input signal may be upsampled by any integer number, e.g., 2, 3, 6 and the like, it is preferred that upsampling be by a factor of 2. This ensures that the signal w_k received by the interpolator 604 has no energy above half the Nyquist frequency associated therewith, which due to the upsampling by a factor of two is twice the Nyquist frequency of the incoming x_n . Effectively, the intermediate signal w_k has a much wider guardband than the incoming signal x_n . For example, for an incoming 48 kHz signal x_n with an original guardband extending from 20 kHz to 24 kHz, the guardband of the upsampled,

filtered, and intermediate signal, w_k at sample rate 96 kHz extends from 20 kHz to 48 kHz. Thus the width of the guardband has been increased from 4 kHz to 28 kHz, a factor of seven. With such a dramatically wider guardband, the transition width of the filter associated with interpolator 304 can be drastically increased and its complexity correspondingly reduced.

- 5 For example, only a fourth order filter is required for 0.01 dB ripple and 96 dB stop-band attenuation, compared to a requirement of 28th order to achieve the same specifications without upsampling. Table 1 shows quality specifications for various orders of the interpolator 604 as follows:

Order	48 kHz Output Rate				44.1 kHz Output Rate			
	GB Aliases OK		No GB Aliases		GB Aliases OK		No GB Aliases	
	Ripple	SB	Ripple	SB	Ripple	SB	Ripple	SB
4	0.01	-96	0.015	-96	0.14	-96	0.3	-96
5	0.001	-112	0.001	-99	0.01	-105	0.013	-96
6	0.0001	-128	0.0001	-113	0.001	-116	0.001	-104
7	0.00001	-144	0.00001	-126	0.0001	-128	0.0001	-114
8	0.000001	-159	0.000002	-152	0.00002	-142	0.0001	-143

TABLE 1

When linear phase is a system requirement, it is preferred that upsampling filter 602 be an FIR half-band upsampling filter. When employed with the interpolator 604, the system's computational complexity can be approximately halved compared to forming a converter 600 without the filter 602, while satisfying the same converter quality requirements. An example of the filter 602 orders required to achieve certain levels of quality for two standard sample rates is shown below in Table 2.

<u>48 kHz Input Rate (Passband edge 5/6 Fn)</u>			<u>44.1 kHz Input Rate (Passband edge 10/11 Fn)</u>		
<u>Order</u>	<u>Ripple</u>	<u>SB atten</u>	<u>Order</u>	<u>Ripple</u>	<u>SB atten</u>
77	0.00015	-96	141	0.00015	-96
89	0.00004	-108	157	0.00004	-108
5 101	0.00001	-120	177	0.00001	-120
113	0.000001	-132	197	0.000002	-132
121	0.0000002	-144	217	0.0000005	-144

TABLE 2

10 It is preferred that the half-band upsampling filter attenuate the first image of the signal x_n . The actual filter design criterion is that the half-band upsampling filter passband extend through the passband of x_n , that the filter transition band extend from the passband edge of x_n to the first image of the passband edge (thus having a transition band width of twice the guardband of x_n and having the stopband extend from the first image of the passband edge to the new Nyquist frequency at the sample rate of x_n). Any number of FIR filter design methods could be used to design this filter, but the preferred method is a windowed sinc function. A Chebyshev window is typically used, but a Kaiser window will produce similar results. Because of the use of a windowed sinc function, the computational complexity of an FIR half-band upsampling filter is reduced. This is because the value of the center coefficient will be unity, and the value of all coefficients spaced an even number away from the center will be zero. Consequently, when so designed, the filter 602 passes the input data unchanged to produce half the output data as unaltered sample points. When the remaining half of the data, which are interpolated points, are computed, the left and right hand lobes of the filter are symmetrical, so the associated data points can be added together before multiplication and summation. Thus, if the order of the half-band upsampling filter 602 is $4n+1$, a total of n multiplicative operations and $2n$ additive operations are performed for each input sample.

25 The computational complexity of the half-band upsampling filter is proportional to the rate of the input, because the filter operates on the incoming sample stream. As a result, the maximum input rate is nominally the output rate. In the case of hardware implementations, reservations may be made such that the input rate can slightly exceed the output rate, so that inaccuracy and drift can be accommodated when receiving asynchronous input data at nominally the same sample rate.

First Embodiment: Halfband filter in upsampler when output sample rate approx. equal to or higher than input sample rate.

The parameters of the upsampling filter and of the interpolator for this case when the output signal has approximately the same or wider passband than the input signal (which is generally the case when the output sample rate is approximately equal to or higher than the input rate) are illustrated in Figs. 7A-7E. Fig. 7A shows the spectrum of the signal X_n , including passband frequency range 701, guardband 702, first image 703, second image 704, and third image 705. Images above the third are not shown.

The desired parameters for half-band upsampling filter 602 are shown in Fig.

7B. These include passband ripple 711, transition band 712 of width less than or equal to twice the guardband whose lower end 717 must be no lower than the upper edge of input signal passband 701, and whose upper end 718 must be below the lower edge of the first passband image 703, and stopband 713 with finite attenuation. Also shown are the first images of the filter stopband, transition band and passband as 714, 715 and 716 respectively, and first image 719 of the lower end 717 of transition band 712.

Fig. 7C shows the spectrum of intermediate signal W_k , including passband 721 and the lowest unfiltered passband image 722.

Fig. 7D shows the desired parameters for the filter of interpolator 604, including passband ripple 731, wider transition band 732 whose lower end is no lower than the lower end 717 of half-band filter transition band 712, and whose upper end is below the first image 719 of the lower end 717 of half-band filter transition band 712, and stopband 733 with finite attenuation.

Fig. 7E shows the result of the interpolator filtering operation with passband 741 unchanged (apart from the effects of passband ripple) and all images attenuated by the stopbands, thus preventing any significant aliasing when decimation occurs in the final interpolator operation.

It should be clear that because the passbands of the half-band upsampling filter and the interpolator are both applied to the signal passband, in general the sum of these parameters must be below the required system ripple specification. In practice it will be found that the passband ripple of half-band filters with appropriate stopband rejection is exceedingly small compared to that of the interpolator filter, and thus this situation is easily arranged. Similarly it should be clear that in the general case, the transition band and stopband regions of the filters ultimately have both filters applied, thus in some instances

advantage may be taken of this situation to relax the constraints while still achieving the required system requirements.

For purposes of comparison, the preferred embodiment requires a 4th order interpolator and a 77th order half-band upsampling filter, for a total complexity of 19+8=27 multiplications and 38+12=50 additions, while the prior art classical algorithm requires a 28th order interpolator of complexity 56 multiplications and 84 additions to achieve the same specifications of 0.01 dB ripple and 96 dB stopband attenuation. An additional benefit of the preferred embodiment over the prior art is a reduction in the table size for the coefficients for the interpolator by a factor of seven.

10 IIR Upsampling Filter where Linear Phase not Required.

Where linear phase is not a system requirement, an IIR half-band upsampling filter offers still further computational advantages, compared with the aforementioned FIR half-band filter. A particular form of IIR filter, the power symmetric elliptical filter, is optimal for this purpose. Such filters can be implemented in a variety of ways known to those skilled in the art. A preferred implementation flow diagram for a sixth order filter is shown in Fig. 8A. Input 801 is passed through first order allpass filters 802, 803, and 804, and in parallel through first order allpass filters 805, 806 and 807. Output 808 commutes between the outputs of allpass 804 and 807 to produce the upsampled signal at twice the input rate. Each first order allpass filter can be constructed in a variety of ways known to those skilled in the art. A preferred implementation flow diagram is shown in Fig. 8B. Allpass input 811 is summed in adder 812 with the output of unit delay element 813 and multiplied by coefficient K in multiplier 814. Subtractor 815 deducts the output of multiplier 814 from input 811, and applies the difference to unit delay element 813. Adder 816 sums the output of unit delay element 813 with the output of multiplier 814 to produce allpass output 817.

25 The values of the coefficients K can be determined according to the references for any required specifications.

Table 3, below, shows the specifications for an exemplary IIR power symmetric elliptical half-band upsampling filter:

48 kHz Input Rate (Passband Edge 5/6 Fn) 44.1 kHz Input Rate (Passband Edge 10/11 Fn)

<u>Order</u>	<u>Ripple</u>	<u>SB atten</u>	<u>Order</u>	<u>Ripple</u>	<u>SB atten</u>
6	<.0000001	-96	6	<.0000001	-78
8	<.0000001	-128	8	<.0000001	-103
10	<.0000001	-159	10	<.0000001	-129
12	<.0000001	-190	12	<.0000001	-155

TABLE 3

10 As can be seen, very substantial gains in computational complexity derive from use of the IIR filter. A power symmetrical elliptical upsampling filter of order n requires n multiplicative operations and 3n additive operations for each input sample. Thus, for example, an implementation at 48 kHz with 0.01 dB ripple and 96 dB stopband rejection takes only a 4th order interpolator, and an order 6 IIR half-band upsampling filter. This requires $6+8 = 14$ multiplicative operations and $18+12 = 30$ additions, reducing multiplicative operations to 25%, and additive operations to 36%, of that required by the classical algorithm. At lower input rates, the efficiency is even higher. The phase non-linearity of the IIR filter can also be reduced by addition of a low order allpass stage with a compensating phase characteristic in series with the IIR filter. Fig. 9 shows such a compensating circuit, where first order allpass filter 901 partially compensates for the phase nonlinearity of power symmetric elliptical upsampling half-band filter 902.

Flow Chart (Fig. 10).

Referring to Fig. 10, in operation, any one of the embodiments discussed above with respect to the sample rate converter 600 would receive a plurality of data points, 25 associated with the signal x_n , at a first sample rate at step 1010. At step 1012, the filter 602 would operate on the plurality of data points to associate the signal with a predetermined set of parameters forming signal w_k . The parameters associated with the signal are stopband rejection, passband ripple, numerical accuracy and transition band having a first width. At step 1014, the sample rate associated with the input signal x_n and, therefore, the intermediate 30 signal w_k is varied by interpolating a subset of data points of the plurality of data points with the interpolator 604. Having associated therewith a transition band with a width that is greater than the width associated with the transition band of the intermediate signal w_k , the interpolator 604 is much less computationally complex, for the reasons discussed above. Specifically, during step 1014, the interpolator 604 varies the sample rate by convolving a

predetermined finite number N of data points of the signal w_k with an equal number of coefficients derived from those stored in ROM 506. Typically, N is determined by the desired converter quality and the allowable width at the second transition band, and the values of the coefficients are typically a function of the temporal spacing between the output 5 data sample and the corresponding data points. In this manner, the sample rate converter increases, or changes, to a different but nominally similar rate, the sample rate of signal x_n .

Second Embodiment: Halfband filter instead of upsampling when input sample rate is approximately twice output sample rate.

Referring to Figs. 5 and 11, in another embodiment of the present invention, 10 information stored in ROM 506 defines a sample rate converter 1100 that converts an incoming stream of data samples x^1_n , representing an input signal, at a first sample rate to a second, lower, sample rate y^1_m . As before, oversampling can again serve to reduce the computational complexity of converter 1100. For example, when the input sample rate is nominally (but not necessarily exactly) twice the output sample rate, a half-band (but not 15 upsampling) filter may be employed with an interpolator. Rather than oversampling the input signal x^1_n by a factor of two, as discussed above, the samples representing the input signal x^1_n are simply filtered. Either FIR or IIR filters may be employed, and the same coefficients derived previously can be used. To that end, the converter 1100 includes a half-band filter 1102 coupled to an interpolator 1104 such that the data samples x^1_n are received by the filter 20 1102. Such a configuration is particularly useful for converting from a sample rate of nominally 96 kHz to a local 48 kHz rate. By placing slightly tighter constraints on the interpolator 1104, the same system can also convert 88.2 kHz sample rate data to the local 48 kHz rate.

The parameters of the half-band filter 1102 and of the interpolator 1104 for 25 this case are illustrated in Figs. 12A-12G. Fig. 12A shows the spectrum of the input signal X^1_n , including passband frequency range 1201, first image 1203, second image 1204, and third image 1205. Images above the third are not shown.

Fig. 12B shows the spectrum allocation of the output signal Y^1_m , including 30 passband frequency range 1211, guardband 1212, first image 1213, second image 1214, and third image 1215. Images above the third are not shown.

The desired parameters for half-band filter 1102 are shown in Fig. 12C. These include passband ripple 1221, transition band 1222 of width less than or equal to twice the output guardband whose lower end 1227 must be no lower than the upper edge of output signal passband 1211, and whose upper end 1228 must be below the lower edge of the first

output passband image 1213, stopband 1223 with finite attenuation, and first images of the stopband, transition band and passband as 1224, 1225 and 1226 respectively. Also shown is the first image 1229 of transition band upper end 1228.

Fig. 12D shows the spectrum of intermediate signal, including passband 1231, partially attenuated input signal 1232, the first image of the partially attenuated signal 1233 and first image of passband 1234.

Fig. 12E shows the desired parameters for the filter of interpolator 1104, including passband ripple 1241, wider transition band 1242 which must begin no lower than the lower end 1227 of half-band filter transition band 1222, and which must end below the first image 1229 of the upper end 1228 of half-band filter transition band 1222, and stopband 1243 with finite attenuation.

Fig. 12F shows the result of the interpolator filtering operation with passband 1251 unchanged (apart from the effects of passband ripple) and all images attenuated by the stopbands except for the partially attenuated input signal 1253.

Fig. 12G shows the output signal Y_m^1 , including passband 1261 unchanged (apart from the effects of passband ripple) and partially attenuated input signal 1263 and its alias 1264 limited entirely to within the output guardband, thus containing no significant aliasing within the passband.

Third Embodiment: Interpolator followed by halfband decimating filter for output sample rate less than or approximately equal to, but more than half, the input sample rate.

Referring to Figs. 5 and 13, were the output sample rate less than or approximately equal to the input sample rate, but well above half of the same, then it would be preferred to have a sample rate converter 1300 with an interpolator 1304 producing an intermediate signal at exactly twice the output sample rate. The interpolator 1304 is followed by a half-band decimating filter 1306 having a sharp cutoff. In this manner, the interpolator 1304 need only reduce aliasing above the bottom of the fourth image of the passband at the output sample rate, because the half-band decimating filter 1306 effectively eliminates everything below this image. In other words, the transition band of interpolator 1304 can extend from the edge of the passband of the (lower rate) output signal through the first image of the passband edge of the input signal. For example, were the converter used to convert 48 kHz sample rate data to an output rate of 44.1 kHz, the interpolator's transition band could extend from 20 kHz to 68.2 kHz. The half-band decimating filter 1306 may be implemented as an FIR or IIR filter, depending on system requirements. Were an IIR filter selected, the

preferred embodiment would include a power symmetric half-band decimating elliptical filter, which can be implemented in a variety of ways known to those skilled in the art.

A preferred implementation flow diagram for a sixth order version of such a power symmetric half-band decimating elliptical filter is shown in Fig. 8C. Input 821 is alternately commutated to a chain of first order allpass filters 822, 823, and 824, and to the other chain of first order allpass filters 825, 826 and 827. Adder 828 sums the outputs of allpass 824 and 827 to produce the decimated output signal 829 at half the input rate. As in the other embodiments, the computational efficiency of the half-band filter in providing a sharp transition band allows the combined system to be substantially more computationally efficient than traditional methods.

The parameters of the half-band decimating filter 1306 and of the interpolator 1304 (shown in Fig. 13) for this example are illustrated in Figs. 14A-14G. Fig. 14A shows the spectrum of the 48 kHz input signal, including passband frequency range 1401, first image 1403, second image 1404, and third image 1405. Images above the third are not shown.

Fig. 14B shows the spectrum allocation of the 44.1 kHz output signal, including passband frequency range 1411, guardband 1412, first image 1413, second image 1414, and third image 1415. Images above the third are not shown.

The desired parameters for half-band decimating filter 1306 are shown lower in Fig. 14E. These include passband ripple 1421, transition band 1422 of width less than or equal to twice the output guardband whose lower end 1427 must be no lower than the upper edge of output signal passband 1411, and whose upper end 1428 must be below the lower edge of the first output passband image 1413, stopband 1423 with finite attenuation, and first images of the stopband, transition band and passband as 1424, 1425 and 1426 respectively.

Also shown is the first image 1429 of transition band lower end 1427.

Fig. 14C shows the desired parameters for the filter of interpolator 1304, including passband ripple 1441, wider transition band 1442 which must begin no lower than the lower end 1427 of half-band filter transition band 1422, and which must end below the first image 1429 of the lower end 1427 of half-band filter transition band 1422, and stopband 1443 with finite attenuation.

Fig. 14D shows the intermediate signal produced as a result of the interpolator filtering operation with passband 1451 unchanged (apart from the effects of passband ripple) and all images attenuated by the stopband except for the partially attenuated input images 1453.

Fig. 14F shows the spectrum of the signal resulting from the half-band filtering operation prior to decimation, including passband 1431 and partially attenuated input image 1432.

Fig. 14G shows the decimated output signal, including passband 1461 unchanged (apart from the effects of passband ripple) and partially attenuated input image alias 1463 limited entirely to within the output guardband, thus containing no significant aliasing within the passband.

Fourth Embodiment: Halfband filter on both upsampling and decimation where dynamically changing sample rate.

Referring to Fig. 15, the advantage of the preferred embodiment of the interpolator in all of the aforementioned sample rate converters is that it facilitates providing a dynamically changing sample rate. As a result, the sample rate ratio may dynamically change. Consider, for example, a sample rate converter 1500 that provides sample rate ratios from slightly more than 1/3 to infinity (in practice an arbitrarily high value). In this case, the input signal x^2_n is first upsampled by a factor of two using a half-band upsampling filter 1502. The sample rate is then converted to twice the desired output rate using a wide transition band interpolator 1504. Finally, the output of the interpolator is decimated by a factor of two using a half-band decimating filter 1506, as described above producing output signal y^2_m .

The parameters of the upsampling and decimating filters and of the interpolator for this case are illustrated in Fig. 16. Fig. 16A shows the spectrum of the signal X^2_n , including passband frequency range 1601, guardband 1602, first image 1603, second image 1604 and third image 1605. Images above the third are not shown.

The desired parameters for half-band upsampling filter 1502 are shown in Fig. 16B. These include passband ripple 1611, transition band 1612 of width less than or equal to twice the guardband 1602 whose lower end 1617 must be no lower than the upper edge of input signal passband 1601, and whose upper end 1618 must be below the lower edge of the first passband image 1603, and stopband 1613 with finite attenuation. Also shown are the first images of the filter stopband, transition band and passband as 1614, 1615 and 1616 respectively, and first image 1619 of the lower end 1617 of transition band 1612.

Fig. 16C shows the spectrum of the first intermediate signal, including passband 1621 and the lowest unfiltered passband image 1622.

Fig. 16D shows the desired parameters for the filter of interpolator 1504, including passband ripple 1631, wider transition band 1632 whose lower end is no lower than

the lower end 1617 of half-band filter transition band 1612, and whose upper end is below the first image 1619 of the lower end 1617 of half-band filter transition band 1612, and stopband 1633 with finite attenuation.

Fig. 16E shows the result of the interpolator filtering operation with passband 5 1641 unchanged (apart from the effects of passband ripple) and all images attenuated by the stopbands.

It will be noted at this point that the operation is identical to that of the embodiment of Fig. 10, except that the interpolator 1504 is producing an intermediate signal of a sample rate twice the output rate which is yet to be applied to half-band decimating filter 10 1506. In the case when the input sample rate is lower or approximately equal to the output rate, the input passband 1601, now available at the interpolator output as passband 1641, will be smaller than the available output passband, and thus not require any further filtering. Indeed, half-band decimating filter 1506 will leave this entire smaller passband unchanged as is desired. But if the input sample rate exceeds the output sample rate, which in this case it can do by nearly a factor of three, the passband 1641 requires further filtering to prevent aliasing on decimation.

Fig. 16F shows the spectrum allocation of the output signal which has been scaled for the interesting case when it is slightly higher than one-third the input sample rate, including passband frequency range 1651, guardband 1652, first image 1653, second image 15 1654, and third image 1655. Images above the third are not shown.

The desired parameters for half-band decimating filter 1506 are shown lower in Fig. 16G. These include passband ripple 1661, transition band 1662 of width less than or equal to twice the output guardband whose lower end 1667 must be no lower than the upper edge of output signal passband 1651, and whose upper end 1668 must be below the lower 20 edge of the first output passband image 1653, stopband 1663 with finite attenuation, and first images of the stopband, transition band and passband as 1664, 1665 and 1666 respectively.

Fig. 16H shows the spectrum of the signal resulting from the final half-band filtering operation prior to decimation, including passband 1671 and partially attenuated intermediate signal 1672.

Fig. 16I shows the decimated output signal, including passband 1681 unchanged (apart from the effects of passband ripple) and partially attenuated intermediate signal 1682 and its alias 1683 limited entirely to within the output guardband, thus containing 30 no significant aliasing within the passband.

While all of the above embodiments presume the less conservative interpretation that allows guardband aliasing to occur, the same principles apply to the design of sample rate converters that take the more conservative approach which precludes guardband aliasing. The scope of the invention, therefore, should not be determined with reference to the above description, but should be determined with reference to the appended claims along with their full scope of equivalents.

WHAT IS CLAIMED IS:

1 1. A method for converting a signal to differing sample rates, comprising:
2 receiving, at a first sample rate, a plurality of data points, associated with a
3 first signal;
4 operating on said plurality of data points to associate said signal with a
5 predetermined set of parameters, with said set of parameters including a first transition band
6 having an image corresponding thereto; and
7 varying said sample rate associated with said first signal by interpolation with
8 an interpolator having associated therewith a second transition band, with the width
9 associated with said second transition band being a function of a spectral separation of said
10 first transition band and said image, wherein a second signal is produced having a sequence
11 of data samples approximating the first signal.

1 2. The method recited in claim 1 wherein varying said sample rate
2 includes producing each data sample associated with said second signal by convolving a
3 predetermined finite number N of data points with an equal number of coefficients, with N
4 being greater than two.

1 3. The method recited in claim 2 wherein coefficients vary as a function
2 of the temporal spacing between the output point and the corresponding input points.

1 4. The method as recited in claim 1 wherein varying said sample rate
2 increases said sample rate.

1 5. The method as recited in claim 1 wherein varying said sample rate
2 decreases said sample rate.

1 6. The method as recited in claim 1 wherein operating on said plurality of
2 data points includes up-sampling said plurality of data points by a factor of two.

1 7. The method as recited in claim 1 wherein operating on said plurality of
2 data points includes filtering said plurality of data points with a half-band filter.

1 8. The method as recited in claim 1 wherein operating on said plurality of
2 data points includes decimating said plurality of data points with a half-band decimator.

1 9. The method as recited in claim 6 further including decimating a
2 plurality of data points output by said interpolator with a half-band decimator, with varying
3 said sample rate occurring after receiving said plurality of data points and before decimating
4 said plurality of data points.

1 10. The method as recited in claim 1 wherein operating on said plurality of
2 data points to associate said signal includes filtering the same with a finite impulse response
3 filter.

1 11. The method as recited in claim 1 wherein operating on said plurality of
2 data points to associate said signal includes filtering the same with an infinite impulse
3 response filter.

1 12. A method for converting a digital audio signal to a different sample
2 rate, comprising:

3 receiving a plurality of data points, associated with an audio signal, at an
4 initial sample rate;
5 halfband filtering said plurality of data points with a halfband filter; and
6 interpolating with an interpolator having independently programmable
7 parameters.

1 13. The method as recited in claim 12 wherein:
2 said halfband filtering is done in conjunction with upsampling said plurality of
3 data points; and
4 said interpolating follows said upsampling and halfband filtering.

1 14. The method as recited in claim 12 wherein:
2 said halfband filtering is done, without upsampling, on said plurality of
3 datapoints; and
4 said interpolating follows said halfband filtering.

1 15. The method as recited in claim 12 wherein:
2 said halfband filtering follows said interpolating.

1 16. The method as recited in claim 12 wherein:
2 said halfband filtering is done in conjunction with upsampling said plurality of

3 data points;

4 said interpolating follows said halfband filtering; and
5 halfband filtering and decimating following said interpolating.

1 17. A computer program product for converting signals to differing sample
2 rates comprising:

3 code for receiving a plurality of data points, associated with a signal, at a first
4 sample rate;

5 code for operating on said plurality of data points to associate said signal with
6 a predetermined set of parameters, with said set of parameters including a first transition band
7 having a first width;

8 code for varying said sample rate associated with said first signal by
9 interpolating a subset of data points of said plurality of data points with an interpolator
10 having associated therewith a second transition band, with the width associated with said
11 second transition band being a function of a spectral separation of said first transition band
12 and said image, wherein a second signal is produced having a sequence of data samples
13 approximating the first signal; and

14 a computer-readable storage medium for storing code.

1 18. The computer program product as recited in claim 17 wherein code for
2 operating on said plurality of data points includes code for up-sampling said plurality of data
3 points by a factor of two.

1 19. The computer program product as recited in claim 17 wherein code for
2 operating on said plurality of data points includes code for filtering said plurality of data
3 points with a half-band filter.

1 20. The computer program product as recited in claim 17 wherein code for
2 operating on said plurality of data points includes code for decimating said plurality of data
3 points with a half-band decimator.

1 21. The computer program product as recited in claim 18 further including
2 code for decimating said plurality of data points with a half-band decimator.

1 22. The computer program product as recited in claim 17 wherein code for
2 operating on said plurality of data points to associate said signal includes code for filtering

3 said data points with a filter selected from the set of filters consisting essentially of a finite
4 impulse response filter and a infinite impulse response filter.

1 23. A computer program product for converting a digital audio signal to a
2 different sample rate, comprising:

3 a computer-readable storage medium for storing code, said code including
4 code for receiving a plurality of data points, associated with an audio signal, at
5 an initial sample rate;

6 code for halfband filtering said plurality of data points with a halfband filter;
7 and

8 code for interpolating with an interpolator having independently
9 programmable parameters.

1 24. The computer program product as recited in claim 23 wherein:

2 said code for halfband filtering is executable in conjunction with code for
3 upsampling said plurality of data points; and

4 said code for interpolating is executable following said upsampling and
5 halfband filtering code.

1 25. The computer program product as recited in claim 23 wherein:

2 said code for halfband filtering is executable, without upsampling code, on
3 said plurality of datapoints; and

4 said code for interpolating is executable following said halfband filtering code.

1 26. The computer program product as recited in claim 23 wherein:

2 said code for halfband filtering is executable following said code for
3 interpolating.

1 27. The computer program product as recited in claim 23 wherein:

2 said code for halfband filtering is executable in conjunction with code for
3 upsampling said plurality of data points;

4 said code for interpolating is executable following said code for halfband
5 filtering; and further comprising:

6 code for halfband filtering and decimating executable following said code for
7 interpolating.

1 28. The method of claims 1 or 12 wherein said interpolator is an FIR Nth
2 order sum of products interpolator with linear interpolation of coefficients.

1 29. The computer program product of claims 17 or 23 wherein said
2 interpolator is an FIR Nth order sum of products interpolator with linear interpolation of
3 coefficients.

1 30. The method of claims 1 or 12 wherein said interpolator has a transition
2 band beginning adjacent the top of a passband and ending adjacent the bottom of a passband
3 image.

1 31. The computer program product of claims 17 or 23 wherein said
2 interpolator has a transition band beginning adjacent the top of a passband and ending
3 adjacent the bottom of a passband image.

1 32. The method of claims 7, 8, 9 or 12 wherein said halfband filter is an
2 IIR filter composed of first order allpass blocks.

1 33. The computer program product of claims 19 or 23 wherein said
2 halfband filter is an IIR filter composed of first order allpass blocks.

3
4
5


SAMPLE RATE CONVERTER HAVING DISTRIBUTED FILTERING

ABSTRACT OF THE DISCLOSURE

A method and a computer program product for sample rate conversion that features distributive or hybrid filtering to reduce unwanted artifacts, such as aliasing and the

- 5 receiving, at a first sample rate, a plurality of data points, associated with a first signal, operating on the plurality of data points to associate the signal with a predetermined set of parameters, with the set of parameters including a first transition band having an image associated therewith; and varying the sample rate associated with the first signal by
- 10 interpolation with an interpolator having associated therewith a second transition band, with the width associated with the second transition band being a function of a spectral separation between the first transition band and its image, wherein a second signal is produced having a sequence of data samples approximating the first signal.

PA 3008731 v11






Fig. 4A (P: 01 Art)




Fig. 4B (P: 01 Art)




Fig. 4C (P: 01 Art)




Fig. 4D (P: 01 Art)





Fig. 4E (P: 01 Art)



Frequency (kHz) →


*Fig. 5*

Fig. 6








Fig. 7E Interpolation result















Fig. 8C







Fig. 9

*Fig. 10**Fig. 11*



✓

09427815-102799



Frequency →




Fig. 13












Fig. 15



DECLARATION AND POWER OF ATTORNEY

As a below named inventor, I declare that:

My residence, post office address and citizenship are as stated below next to my name; I believe I am the original, first and sole inventor of the subject matter which is claimed and for which a patent is sought on the invention entitled: **SAMPLE RATE CONVERTER HAVING DISTRIBUTED FILTERING** the specification of which is attached hereto.

I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above. I acknowledge the duty to disclose information which is material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56. I claim no foreign priority benefits under Title 35, United States Code, Section 119.

POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorneys and/or agents to prosecute this application and transact all business in the Patent and Trademark Office connected therewith.

Kenneth C. Brooks, Reg. No. 38,393

Paul C. Haughey, Reg. No. 31,836

Send Correspondence to: Paul C. Haughey TOWNSEND and TOWNSEND and CREW LLP Two Embarcadero Center, 8th Floor San Francisco, California 94111-3834		Direct Telephone Calls to: (Name, Reg. No., Telephone No.) Name: Paul C. Haughey Reg. No.: 31,836 Telephone: 650-326-2400	
Full Name of Inventor 1: ROSSUM	Last Name: ROSSUM	First Name: DAVID	Middle Name or Initial: P.
Residence & Citizenship: Monterey	City: Monterey	State/Foreign Country: California	Country of Citizenship: United States
Post Office Address: 25513 Paseo de Cumbre	City: Monterey	State/Country: California	Postal Code: 93940

I further declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application or any patent issuing thereon.

Signature of Inventor 1



Date 13 October 1999